To make Medium work, we log user data. By using Medium, you agree to our Privacy Policy, including cookie policy.

Privacy as a Creative Constraint: When Anonymization Becomes an Aesthetic Language

Matteo Sipione Follow 7 min read · 10 hours ago

10

Privacy constraints are often framed as limitations that reduce what can be

shown, shared, or remembered. In practice, they frequently act as creative pressure. When documentation of real events collides with the need to protect personal identity, especially in visual media, the result is usually subtraction: faces blurred, images discarded, archives locked away. What remains is compliant, but often visually sterile and narratively weakened.

In this case, the starting point was a simple and familiar problem. I had photographs of real people in real spaces, images that carried aesthetic and documentary value, but could not be published or processed further without risking personal identification. Faces were heavily blurred, identities already erased in a legal sense, yet the presence of bodies still raised questions. Instead of pushing the images further into obscurity, I chose to move in the opposite direction and lean into transformation.

By using artificial intelligence on already anonymized photographs, I introduced an additional layer of abstraction. The AI removes the human body entirely while preserving clothing, posture, and spatial context. What remains resembles mannequins, empty shells, or garments suspended in space. The result is neither surveillance nor concealment, but a visual reinterpretation of presence. Privacy is not treated as a constraint to work around, but as the conceptual engine that generates a new artistic language.

THE ANONYMIZATION

Before any image was introduced into an artificial intelligence pipeline, I began with a strict and deliberate anonymization phase. This step was not treated as a formality, but as a foundational requirement. I processed all photographs to make faces completely unrecognizable, removing any realistic possibility of personal identification. The goal was to ensure that, even before any artistic transformation, the material was already fully

compliant with privacy principles.

The anonymization itself was performed programmatically using custom Python code that I designed to detect and obscure faces in a consistent and aggressive manner. This automated pass ensured scalability and uniformity across large image sets. However, I did not consider automation alone sufficient. I then reviewed each image manually, one by one, to verify that no identifiable features remained and that edge cases were handled correctly. Only images that passed this second human validation stage were admitted to the next phase.

```
import os
import glob
import cv2

# ====== CONFIG ======
INPUT_DIR = "./FOLDER"
```

```
OUTPUT_DIR = "./FOLDER_blurred"
MODEL_PATH = "./face_detection_yunet_2023mar.onnx"

IMAGE_EXTS = (".jpg", ".jpeg", ".png", ".bmp", ".tif", ".tiff", ".webp")

# Detection (if too many / too few faces are detected, tweak score_threshold)
SCORE_THRESHOLD = 0.80
NMS_THRESHOLD = 0.30
TOP_K = 500

# ====== HEAVY BLUR (very aggressive) ======
BLUR_SCALE = 2.5   # increase for stronger blur (e.g. 3.0)
MAX_K = 401        # kernel size cap (very large)
REPEATS = 4        # blur repetitions (increase for more "censorship")

def odd(n: int) -> int:
    return n if n % 2 == 1 else n + 1

def list_images(folder: str):
    files = []
    for ext in IMAGE_EXTS:
        files += glob.glob(os.path.join(folder, f"*{ext}"))
        files += glob.glob(os.path.join(folder, f"*{ext.upper()}"))
    return sorted(files)

def heavy_blur_roi(roi_bgr):
    """
    Very aggressive and stable blur:
    - kernel proportional to face size (min ROI dimension)
    - repeated blur passes
    """
    if roi_bgr is None or roi_bgr.size == 0:
        return roi_bgr

    base = max(25, int(min(roi_bgr.shape[:2]) * BLUR_SCALE))
    k = min(MAX_K, odd(base * 2 + 1))  # very large, odd kernel size

    blurred = roi_bgr
    for _ in range(REPEATS):
        blurred = cv2.GaussianBlur(blurred, (k, k), 0)

    return blurred

def main():
    if not os.path.exists(MODEL_PATH):
        raise FileNotFoundError(
            f"YuNet model not found: {MODEL_PATH}\n"
            "Download it from:\n"
            "https://github.com/opencv/opencv_zoo/raw/main/models/face_detectio
```

```
    )

os.makedirs(OUTPUT_DIR, exist_ok=True)

detector = cv2.FaceDetectorYN.create(
    MODEL_PATH,
    "",
    (640, 640),
    score_threshold=SCORE_THRESHOLD,
    nms_threshold=NMS_THRESHOLD,
    top_k=TOP_K,
)

files = list_images(INPUT_DIR)
if not files:
    print(f"No images found in {INPUT_DIR}")
    return

total_faces = 0

for path in files:
    img = cv2.imread(path)
    if img is None:
        print(f"[SKIP] Unable to read: {path}")
        continue

    h, w = img.shape[:2]
    detector.setInputSize((w, h))

    _, faces = detector.detect(img)

    out = img.copy()
    n_faces = 0

    if faces is not None:
        for f in faces:
            x, y, bw, bh = map(int, f[:4])

            # Clamp bounding box
            x1 = max(0, x)
            y1 = max(0, y)
            x2 = min(w, x + bw)
            y2 = min(h, y + bh)

            if x2 <= x1 or y2 <= y1:
                continue

            roi = out[y1:y2, x1:x2]
            out[y1:y2, x1:x2] = heavy_blur_roi(roi)
```

```
                n_faces += 1

        out_path = os.path.join(OUTPUT_DIR, os.path.basename(path))
        ok = cv2.imwrite(out_path, out)
        if not ok:
            print(f"[ERR ] Unable to save: {out_path}")
            continue

        total_faces += n_faces
        print(f"[OK  ] {os.path.basename(path)} -> faces blurred: {n_faces}")

    print(f"Done. Output: {OUTPUT_DIR}. Total faces blurred: {total_faces}")

if __name__ == "__main__":
    main()
```

By the time the AI was involved, the photographs no longer represented people in a legal or biometric sense. They were already abstracted, stripped of identity, and safe to process. This sequencing is critical to the project's logic. I did not use artificial intelligence to anonymize reality after the fact,

but to reinterpret material that had already been ethically and legally neutralized. The creative process therefore starts from a position of compliance, not remediation, and builds meaning on top of that constraint rather than attempting to bypass it.

THE PROMPT

At this point, the anonymized images were ready to be transformed rather than corrected. I deliberately chose to rely on a general-purpose generative system rather than a bespoke vision model, precisely to test whether a mainstream AI could be guided into producing a coherent and controlled visual outcome.

Get Matteo Sipione's stories in your inbox

Join Medium for free to get updates from this writer.

Enter your email

Subscribe

The interaction with the model was entirely prompt-driven. No additional masks, reference images, or post-processing tricks were used. The intent was to describe, with extreme precision, what needed to disappear and what needed to remain untouched. In particular, the prompt enforces a strict separation between the human body and the clothing, treating the body as a removable volume while considering garments as inviolable objects that must stay physically and visually consistent.

The prompt was written as a technical instruction rather than a creative suggestion, leaving as little interpretative freedom as possible. Below is the exact prompt used. It is included verbatim, as code, to make the process

transparent and reproducible.

```
Remove all humans completely from the image.

The human body must be removed as a continuous invisible volume inside the clot

Under no circumstances remove, cut, erase, or replace any clothing items.
All garments — including pants, trousers, jeans, skirts, dresses, tights, stock

Clothing must always remain visible in its entirety.
Pants and lower-body garments must never be removed, shortened, faded, or confu

The absence of the body must be represented only as empty internal space inside
Garments must appear hollow and empty inside, with physically correct interior

Garments with deep necklines, open cuts, or tight fits must show empty interior

Clothes must remain naturally shaped and positioned, with correct gravity, volu

Do not leave any traces of human anatomy, skin, hair, or biological elements.
Do not replace people with mannequins, ghosts, silhouettes, transparent bodies,
Do not alter the background, environment, or non-human objects.
```

This step is where the project fully shifts from compliance to interpretation. The AI is not asked to hide identities or blur reality further, but to construct a new visual state that could not exist without the initial privacy constraint. The result is an image that still documents a scene, yet no longer depicts people as biological subjects, only as the negative space they once occupied.

CONCLUSION

What emerges from this process is not a workaround, but a reframing. I usually see privacy addressed by reducing fidelity, adding noise, or removing content until the image becomes legally acceptable and visually inert. Here, the opposite happens. Compliance becomes the starting condition for creation, not its limitation. By fully anonymizing first and transforming later, I avoid the ambiguity of retroactive justification and instead build an aesthetic language that could only exist because privacy

was taken seriously.

The resulting images are neither portraits nor erasures. They preserve context, posture, social space, and material culture while removing the biological subject entirely. Clothing becomes a trace of presence rather than an attribute of identity. Bodies are not censored; they are conceptually subtracted. What remains is a documentation of absence that is paradoxically more expressive than a blurred face ever could be.

In this sense, I do not use artificial intelligence to extract more information, but to enforce a boundary. I guide the system to respect that boundary with precision, turning a legal and ethical requirement into a formal constraint. The work sits in the space between documentation and fiction, between compliance and interpretation, suggesting that privacy-aware media does not have to be visually impoverished. On the contrary, when treated as a design principle rather than a legal afterthought, privacy can generate entirely new ways of seeing.

AI Art Privacy Gdpr Coding

Written by Matteo Sipione

1 follower · 2 following

I'm just out there doing stuff. 🎥📽️

Follow

No responses yet

Write a response

What are your thoughts?

Recommended from Medium

In AI Advances by Gautam Bharadwaj

Generate Unlimited AI Images and Videos Locally—No Internet, No...

A practical guide to local AI image and video generation

3d ago 222 3

In Generative AI by Adham Khaled

Stanford Just Killed Prompt Engineering With 8 Words (And I...

ChatGPT keeps giving you the same boring response? This new technique unlocks 2×...

Oct 20, 2025 22K 566

InkOfCode

Why Most Engineers Learn Rust & Go Wrong (And What Senior...

The GitHub Graveyard

3d ago 3

In UX Collective by Theresa-Marie Rhyne

A Nano Banana with color theory

Using Google Gemini to build an Accessible Perceptual Uniform Triad

6d ago 186 2

Will Lockett

The AI Bubble Is About To Burst, But The Next Bubble Is Already...

Techbros are preparing their latest bandwagon.

Sep 14, 2025 20K 832

In Level Up Coding by Vivedha Elango

Your AI Agent Is Failing Because of Context, Not the Model

The complete Technical guide to Context Engineering, context selection, memory...

6d ago 363 6

See more recommendations

Help Status About Careers Press Blog Privacy Rules Terms Text to speech